

THE TRIVIALITY PROBLEM FOR PROFINITE COMPLETIONS

MARTIN R. BRIDSON AND HENRY WILTON

ABSTRACT. We prove that there is no algorithm that can determine whether or not a finitely presented group has a non-trivial finite quotient; indeed, this property remains undecidable among the fundamental groups of compact, non-positively curved square complexes. We deduce that many other properties of groups are undecidable. For hyperbolic groups, there cannot exist algorithms to determine largeness, the existence of a linear representation with infinite image (over any infinite field), or the rank of the profinite completion.

1. INTRODUCTION

The basic decision problems for finitely presented groups provided a guiding theme for combinatorial and geometric group theory throughout the twentieth century. Activity in the first half of the century was framed by Dehn's articulation of the core problems in 1911 [17], and it reached a climax in 1957-58 with the proof by Novikov [29] and Boone [9] that there exist finitely presented groups with unsolvable word problem. In the wake of this, many other questions about general finitely presented groups were proved to be algorithmically unsolvable (cf. Adyan [1, 2], Rabin [33], Baumslag–Boone–Neumann [5]). In the decades that followed, the study of decision problems shifted towards more refined questions concerning the existence of algorithms within specific classes of groups, and to connections with geometry and topology. However, certain basic decision problems about general finitely presented groups were not covered by the techniques developed in mid-century and did not succumb to the geometric techniques developed in the 1990s. The most obvious of these is the following: can one decide whether or not a group has a proper subgroup of finite index?

Our main purpose here is to settle this question.

Date: 15 December 2014.

2010 *Mathematics Subject Classification.* 20F10, 20F67, 57M07, (20E18, 20F65).

This work was supported by Fellowships from the EPSRC (both authors) and by a Wolfson Research Merit Award from the Royal Society (first author).

Theorem A. *There is no algorithm that can determine whether or not a finitely presented group has a proper subgroup of finite index.*

The technical meaning of this theorem is that there is a recursive sequence of finitely presented groups G_n with the property that the set of natural numbers

$$\{n \in \mathbb{N} \mid \exists H \subsetneq G_n, |G_n/H| < \infty\}$$

is recursively enumerable but not recursive. More colloquially, it says that the problem of determining the existence of a proper subgroup of finite index is *undecidable*.

We shall strengthen Theorem A by proving that the existence of such subgroups remains undecidable in classes of groups where other basic decision problems of group theory are decidable, such as biautomatic groups and the fundamental groups of compact, non-positively curved square complexes. We include this last refinement in the following geometric strengthening of Theorem A.

Theorem B. *There is no algorithm that can determine if a compact square complex of non-positive curvature has a non-trivial, connected, finite-sheeted covering.*

There are various other natural reformulations of Theorem A (and its refinements), each creating a different emphasis. For emphasis alone, one could rephrase Theorem A as “*the triviality problem for profinite completions of finitely presented groups is undecidable*”: there is no algorithm that, given a finitely presented group G , can decide whether the profinite completion \hat{G} is trivial. More substantially, since all finite groups are linear (over any field) and linear groups are residually finite, we can rephrase our main result as follows:

There is no algorithm that can determine whether or not a finitely presented group has a non-trivial finite-dimensional linear representation (over any field); indeed the existence of such a representation is undecidable even for the fundamental groups of compact, non-positively curved square complexes.

In Section 2 we shall explain how classical work of Slobodskoi [36] on the universal theory of finite groups can be interpreted as a profinite analogue of the Novikov–Boone theorem: by definition, the profinite completion \hat{G} is the inverse limit of the finite quotients of G , and the kernel of the natural homomorphism $G \rightarrow \hat{G}$ consists of precisely those $g \in G$ that have trivial image in every finite quotient of G ; implicitly, Slobodskoi constructs a finitely presented group G in which there is no algorithm to determine which words in the generators represent such

$g \in G$. In the setting of discrete groups, one can parlay the undecidability of the word problem for a specific group into the undecidability of the triviality problem for finitely presented groups by performing a sequence of HNN extensions and amalgamated free products, as described in Section 3. Although the profinite setting is more subtle and does not allow such a direct translation, we will attack the triviality problem from a similar angle, deducing Theorem A from Slobodskoi's construction and the following *Encoding Theorem*. This is the key technical result in this paper; its proof is significantly more complex than that of the corresponding theorem for discrete groups and the details are much harder.

Theorem C (Encoding Theorem). *There is an algorithm that takes as input a finite presentation $\langle A \mid R \rangle$ for a group G and a word $w \in F(A)$ and outputs a presentation for a finitely presented group G_w such that*

$$\hat{G}_w \cong 1 \Leftrightarrow w =_{\hat{G}} 1 .$$

Theorems A and C imply that various other properties of finitely presented groups cannot be determined algorithmically. The properties that we shall focus on, beginning with the property $\hat{G} \cong 1$ itself, are neither Markov nor co-Markov, so their undecidability cannot be established using the Adyan–Rabin method.

Some of the most profound work in group theory in recent decades concerns the logical complexity of (word-)hyperbolic groups. In that context, one finds undecidability phenomena associated to finitely generated subgroups but the logical complexity of hyperbolic groups themselves is strikingly constrained (see, for instance, [35] and [25]). Nevertheless, we *conjecture* that there does not exist an algorithm that can determine if a hyperbolic group has a non-trivial finite quotient (Conjecture 9.5). This conjecture would be false if hyperbolic groups were all residually finite. Indeed, we shall prove (Theorem 9.6) that this conjecture is equivalent to the assertion that there exist hyperbolic groups that are not residually finite.

We shall also prove that, as it stands, Theorem A allows one to establish various new undecidability phenomena for hyperbolic groups. We recall some definitions. The *first betti number* $b_1(\Gamma)$ of a group Γ is the dimension of $H_1(\Gamma, \mathbb{Q})$ and the *virtual first betti number* $vb_1(\Gamma)$ is the (possibly infinite) supremum of $b_1(K)$ over all subgroups K of finite index in Γ . A group is *large* if it has a subgroup of finite index that maps onto a non-abelian free group. Note that if Γ is large then $vb_1(\Gamma) = \infty$.

The following theorem summarizes our undecidability results for hyperbolic groups.

Theorem D. *There do not exist algorithms that, given a finite presentation of a torsion-free hyperbolic group Γ , can determine:*

- (1) *whether or not Γ is large;*
- (2) *for any $1 \leq d \leq \infty$, whether or not $vb_1(\Gamma) \geq d$;*
- (3) *whether or not every finite-dimensional linear representation of Γ has finite image;*
- (4) *for a fixed infinite field k , whether or not every finite-dimensional representation of Γ over k has finite image;*
- (5) *whether or not, for any fixed $d_0 > 2$, the profinite completion of Γ can be generated (topologically) by a set of cardinality less than d_0 .*

Items (1) and (2) are contained in Theorem 9.2, items (3) and (4) are contained in Theorem 9.4, and item (5) is contained in Theorem 8.3.

We shall prove in Section 8 that the profinite-rank problem described in item (5) remains undecidable among residually-finite hyperbolic groups. In that context, the bound $d_0 > 2$ is optimal, because the profinite rank of a residually finite group Γ is less than 2 if and only if Γ is cyclic, and it is easy to determine if a hyperbolic group is cyclic. Furthermore, Theorem 9.6 tells us that for $d_0 \leq 2$, problem (5) is decidable if and only if every hyperbolic group is residually finite.

Item (1) should be contrasted with the fact that there *does* exist an algorithm that can determine whether or not a finitely presented group maps onto a non-abelian free group: this is a consequence of Makanin's deep work on equations in free groups [27].

Our final application is to the isomorphism problem for the profinite completions of groups. The arguments required to deduce this from Theorem A are lengthy and somewhat technical, so we shall present them elsewhere [11].

Theorem E. *There are two recursive sequences of finite presentations for residually finite groups A_n and B_n together with monomorphisms $f_n : A_n \rightarrow B_n$ such that:*

- (1) *$\hat{A}_n \cong \hat{B}_n$ if and only if the induced map on profinite completions \hat{f}_n is an isomorphism; and*
- (2) *the set $\{n \in \mathbb{N} \mid \hat{A}_n \not\cong \hat{B}_n\}$ is recursively enumerable but not recursive.*

This paper is organised as follows. In Section 2 we explain what we need from Slobodskoi's work. In Section 3 we lay out our strategy for

proving Theorem C, establishing the notation to be used in subsequent sections and, more importantly, providing the reader with an overview that should sustain them through the technical arguments in Sections 4 and 5. Theorem C is proved in Section 6 and, with Slobodskoi's construction in hand, Theorem A follows immediately. Sections 5 and 6 form the technical heart of the paper. Many of the arguments in these sections concern malnormality for subgroups of virtually free groups. The techniques here are largely topological, involving the careful construction of coverings of graphs (and, implicitly, graphs of finite groups) and the analysis of fibre products in the spirit of John Stallings [37].

In Section 7 we prove that the existence of finite quotients remains undecidable in the class of non-positively curved square complexes. Section 8 deals with profinite rank, and the remaining results about hyperbolic groups are proved in Section 9.

Acknowledgements. We first tried to prove Theorem A at the urging of Peter Cameron, who was interested in its implications for problems in combinatorics [15, 12]; we are grateful to him for this impetus. We thank Jack Button and Chuck Miller for stimulating conversations about Theorem A and its consequences. Finally, we are grateful for the insightful comments of the anonymous referee.

2. SLOBODSKOI'S THEOREM

In this section we explain how the following theorem is contained in Slobodskoi's work on the universal theory of finite groups [36]. We write $F(A)$ to denote the free group on a set A .

Theorem 2.1. *There exists a finitely presented group $G \cong \langle A \mid R \rangle$ in which there is no algorithm to decide which elements have trivial image in every finite quotient. More precisely, the set of reduced words*

$$\{w \in F(A) \mid w \neq_{\widehat{G}} 1\}$$

is recursively enumerable but not recursive.

The theorem that Slobodskoi actually states in [36] is the following.

Theorem 2.2 ([36]). *The universal theory of finite groups is undecidable.*

Slobodskoi's proof of Theorem 2.2 is clear and explicit. It revolves around a finitely presented group $G = \langle a_1, \dots, a_n \mid r_1, \dots, r_m \rangle$ that encodes the workings of a 2-tape Minsky machine M that computes a partially recursive function. Associated to this machine one has a disjoint pair of subsets $S_0, S_1 \subseteq \mathbb{N}$ (denoted X and Y in [36]) that

are *recursively inseparable*: S_0 is the set of natural numbers k such that M halts on input 2^k and S_1 is the set of k such that on input 2^k the machine M visits the leftmost square of at least one of its tapes infinitely often. To say that they are recursively inseparable means that there does not exist a recursive set $D \subseteq \mathbb{N}$ such that $S_1 \subseteq D$ and $S_0 \cap D = \emptyset$.

By means of a simple recursive rule, Slobodskoi defines two sequences of words $w_1^{(k)}, w_2^{(k)}$ ($k \in \mathbb{N}$) in the letters $A^{\pm 1}$. He then considers the following sentences in the first-order logic of groups.

$$\Psi(k) \equiv \forall a_1, \dots, a_n [(r_1 \neq 1) \vee \dots \vee (r_m \neq 1) \vee (w_1^{(k)} = w_2^{(k)} = 1)]$$

Note that the sentence $\Psi(k)$ is *false* in a group Γ if and only if there is a homomorphism $\phi : G \rightarrow \Gamma$ such that at least one of $\phi(w_1^{(k)})$ or $\phi(w_2^{(k)})$ is non-trivial. In particular, $\Psi(k)$ is false in some finite group Γ if and only if either $w_1^{(k)} \neq_{\hat{G}} 1$ or $w_2^{(k)} \neq_{\hat{G}} 1$.

Slobodskoi proves that if $k \in S_1$ then $\Psi(k)$ is true in every periodic group (in particular every finite group) [36, Lemma 6]. He then proves that if $k \in S_0$ then $\Psi(k)$ is false in some finite group [36, Lemma 7].

Proof of Theorem 2.1. Let $G = \langle A \mid R \rangle$ be the group constructed by Slobodskoi. The set $\{w \in F(A) \mid w \neq_{\hat{G}} 1\}$ is recursively enumerable: a naive search will eventually find a finite quotient of G in which w survives, if one exists. If the complement $\{w \in F(A) \mid w =_{\hat{G}} 1\}$ were recursively enumerable, then the set

$$D = \{k \in \mathbb{N} \mid w_1^{(k)} =_{\hat{G}} w_2^{(k)} =_{\hat{G}} 1\}$$

would be recursive. But $S_1 \subseteq D$ and $S_0 \subseteq \mathbb{N} \setminus D$, so this would contradict the fact that S_0 and S_1 are recursively inseparable. \square

Remark 2.3. Kharlampovich proved an analogue of Slobodskoi's theorem for the class of finite nilpotent groups [24].

Remark 2.4. It follows easily from Theorem 2.1 and the Hopfian property of finitely generated profinite groups that there does not exist an algorithm that, given two finite presentations, can determine if the profinite completions of the groups presented are isomorphic or not. It is much harder to prove that the isomorphism problem remains unsolvable if one restricts to completions of finitely presented, residually finite groups [11].

3. A STRATEGY FOR PROVING THEOREM C

In this section we lay out a strategy for proving our main technical result, Theorem C. It is useful to think of Theorem C as a machine

that, given a word in the seed group G , produces a group G_w so that the (non)triviality of $w \in \hat{G}$ is translated into the (non)triviality of the profinite completion \hat{G}_w . Although the techniques required to prove this are quite different from the arguments used to prove the corresponding result for discrete groups (which are straightforward from a modern perspective), the broad outline of the proof in that setting will serve us well as a framework on which to hang various technical results. The notation established here will be used consistently in later sections.

3.1. The discrete case. We fix a finitely presented group $G = \langle A \mid R \rangle$ and seek an algorithm that, given a word $w \in F(A)$, will produce a finitely presented group G^w so that $G^w \cong \{1\}$ if $w =_G 1$ and $G \hookrightarrow G^w$ if $w \neq_G 1$. The first such algorithm was described by Adyan [1, 2] and Rabin [33]. There are many ways to vary the construction; cf. [20].

Replacing G by $G * \langle a_0 \rangle$ and $a \in A$ by $a' = aa_0$, if necessary, we may assume that $A = \{a_0, \dots, a_m\}$ where each a_i has infinite order. And replacing w by $[w, a_0]$, we may assume that if w is non-trivial in G then it has infinite order.

Let $G_1 = G * \langle b_0, \dots, b_m \rangle / \langle\langle w^{b_i} = a_i \mid i = 0, \dots, m \rangle\rangle$ and let $G_2 = G_1 * \langle b_{m+1} \rangle$. Note that if $w = 1$ then G_2 is freely generated by the b_i , whereas if $w \neq 1$ then G_2 is a multiple HNN extension of G with stable letters b_i , whence the natural map $G \rightarrow G_2$ is injective. Choose $m + 2$ words that freely generate a subgroup of the normal closure of $w \in F(w, b_{m+1})$, say $c_j = (w^{b_{m+1}})^{j+1} w (w^{b_{m+1}})^{-1-j}$. Define $F_0 := \langle b_0, \dots, b_m \rangle < G_1$, and further define subgroups of G_2 by

$$F_1 := \langle b_0, \dots, b_{m+1} \rangle \quad F_2 := \langle c_0, \dots, c_{m+1} \rangle \quad F := \langle F_1, F_2 \rangle.$$

The subgroup $F_1 < G_2$ is free of rank $m + 2$. If $w =_G 1$ then $F_2 < G_2$ is trivial. If $w \neq_G 1$ then F_2 is free of rank $m + 2$ and $F = F_1 * F_2$ is the free product.

We take two copies G_2 and G'_2 of G_2 and distinguish the elements and subgroups of G'_2 by primes. Define G^w to be the quotient of $G_2 * G'_2$ by the relations

$$\{c_i = b'_i, b_i = c'_i \mid i = 0, \dots, m + 1\}.$$

If $w =_G 1$ then $G^w \cong 1$. If $w \neq_G 1$ then G^w is an amalgamated product

$$G_2 *_F G'_2$$

where the isomorphism $F \cong F'$ identifies F_1 with F'_2 and F_2 with F'_1 . In particular, the natural map $G \rightarrow G_2 \rightarrow G^w$ is injective and $G^w \not\cong 1$.

3.2. The profinite case. Given $G = \langle A \mid R \rangle$ and $w \in F(A)$, we have to construct, in an algorithmic manner, a finite presentation for a group G_w so that $\widehat{G}_w = 1$ if and only if $w =_{\widehat{G}} 1$. The difficult thing to arrange is that G_w must have some non-trivial finite quotient if $w \neq_{\widehat{G}} 1$.

Remark 3.1. Of the many problems one faces in adapting the preceding argument to the profinite setting, the most fundamental concerns our use of HNN (equivalently, Bass–Serre) theory to see that the natural map $G \rightarrow G^w$ is injective if $w \neq 1$. Sobering examples in this connection are the *simple* groups of Burger and Mozes [13]: these are amalgamations $L_1 *_{\Lambda_1 \cong \Lambda_2} L_2$ where $L_1 \cong L_2$ is a finitely generated free group and $\Lambda_i < L_i$ is a subgroup of finite index. (Earlier examples in a similar vein were given by Bhattacharjee [8] and Wise [43].)

Step 1: controlling the order of the generators a_i and of w . What matters now is the order of a_i and w in finite quotients of G . In order to retain enough finite quotients after performing the HNN extensions in step 2, we must ensure that if $w \neq 1$ in \widehat{G} then w and the generators all have the same order in *some* finite quotient of G (or a proxy of G). It will transpire that in fact we need significantly more control than this. This control is established in Section 4, where the key result is Theorem 4.3.

Step 2: a map $G \rightarrow \widehat{G}_1$ whose image is trivial iff $w = 1$. We define G_1 as above, making the a_i conjugate to w . If $w =_{\widehat{G}} 1$ then $G \rightarrow \widehat{G}_1$ is trivial and \widehat{G}_1 is the profinite completion of the free group F_0 on the stable letters b_i . When $w \neq_{\widehat{G}} 1$, we obtain finite quotients of G_1 in which w survives. But this is not enough: for reasons that will become apparent in step 4, we have to work hard to find virtually free quotients $\eta : G_1 \rightarrow \Gamma_0$ where F_0 injects and is *malnormal*.

We remind the reader that a subgroup $H < G$ is termed *malnormal* if $g^{-1}Hg \cap H = 1$ for all $g \notin H$. This is the central concept of Section 5 and continues to be a major focus in Section 6.

Step 3: the construction of Γ and F . In $G_2 = G_1 * \langle b_{m+1} \rangle$ we have to demand far more of the subgroup F_2 than in the discrete case. Consequently, a much more subtle construction of the elements c_i is required, and this is the subject of Section 5. If $w =_{\widehat{G}} 1$ then F_2 is trivial in every finite quotient of G_2 . If $w \neq_{\widehat{G}} 1$ then $F_1 \cong F_2$ and $F \cong F_1 * F_2$ injects into a virtually free quotient Γ of G_2 (Lemma 6.6) where it is *malnormal* (Proposition 6.9).

Step 4. With our more sophisticated definition of c_i and F in hand, we define G_w to be the quotient of $G_2 * G'_2$ by the relations

$$\{c_i = b'_i, b_i = c'_i \mid i = 0, \dots, m+1\}.$$

It is clear that $\hat{G}_w = 1$ if $w =_{\hat{G}} 1$. If $w \neq_{\hat{G}} 1$, then G_w maps onto $\Gamma *_{F \cong F'} \Gamma'$; as a malnormal amalgamation of virtually free groups, this is residually finite, by a theorem of Wise [42].

Remarks 3.2. (1) A crucial feature of the above process is that each step is algorithmic: judicious choices are made, but these choices depend in an algorithmic manner on the parameter w alone. In particular, the algorithm gives an explicit finite presentation for G_w .

(2) The definition of G_w makes no assumption about the existence or nature of the finite quotients of G in which w has non-trivial image. Equally, the proof that G_w has a non-trivial finite quotient if $w \neq_{\hat{G}} 1$ requires only the *existence* of a finite quotient in which w has non-trivial image; it does not require any knowledge about the nature of such a quotient.

4. A STRENGTHENING OF OMNIPOTENCE

The main result of this section (Theorem 4.3) strengthens Wise's theorem on the *omnipotence* of free groups [41].

Given a virtually free group Γ and a finite list of elements $\gamma_1, \dots, \gamma_n \in \Gamma$, we would like to control the (relative) orders of these elements in finite quotients of Γ . Ideally, we would like to dictate orders arbitrarily, but this is too much to expect. For example, if γ_1 and γ_2 have conjugate powers in Γ , then the possible orders for the image of γ_2 are constrained by those of γ_1 . To isolate this problem, we make the following definition.

Definition 4.1. Let Γ be a group. Elements $\gamma_1, \gamma_2 \in \Gamma$ of infinite order are said to be *independent* if no non-zero power of γ_1 is conjugate to a non-zero power of γ_2 . An m -tuple $(\gamma_1, \dots, \gamma_m)$ of elements from Γ is *independent* if γ_i and γ_j are independent whenever $1 \leq i < j \leq m$.

The next definition makes precise the idea that the orders of independent sets of elements can be controlled in finite quotients.

Definition 4.2. A group Γ is *omnipotent* if, for every $m \geq 2$ and every independent m -tuple $(\gamma_1, \dots, \gamma_m)$ of elements in Γ , there exists a positive integer κ such that, for every m -tuple of natural numbers (e_1, \dots, e_m) there is a homomorphism to a finite group

$$q : \Gamma \rightarrow Q$$

such that $o(q(\gamma_i)) = \kappa e_i$ for $i = 1, \dots, m$.

The preceding definitions are due to Wise [41], who proved that free groups are omnipotent. Bajpai extended this to surface groups [4], and the second author proved that all Fuchsian groups are omnipotent [38]. It follows from Wise's recent deep work on special cube complexes (specifically, from the Malnormal Special Quotient Theorem [40]), that virtually special groups are omnipotent. In particular, virtually free groups are known to be omnipotent. However, we do not want to obscure our current setting with the extra complications of special cube complexes and, more importantly, Wise's method of proof does not provide the additional strengthening contained in item (2) of the following theorem. This refinement is a vital component of the strategy described in the previous section: it will be needed to establish malnormality in Lemma 6.4 and Proposition 6.9.

Theorem 4.3. *Let Γ be a virtually free group and let $(\gamma_1, \dots, \gamma_m)$ be an independent m -tuple of elements of Γ . There is a positive integer κ such that, for every m -tuple of positive integers (e_1, \dots, e_m) , there is a homomorphism to a finite group*

$$q : \Gamma \rightarrow Q$$

such that:

- (1) $o(q(\gamma_i)) = \kappa e_i$ for $i = 1, \dots, m$; and,
- (2) furthermore, $\langle q(\gamma_i) \rangle \cap \langle q(\gamma_j) \rangle = 1$ whenever $i \neq j$.

The following lemma is a key step in the proof of omnipotence for free groups [41] (see also [38, Proposition 4.1]).

Lemma 4.4. *Let Λ be a finitely generated free group. If $(\gamma_1, \dots, \gamma_m)$ is an independent m -tuple in Λ , then there exists a subgroup $F < \Lambda$ of finite index and homomorphisms $\phi_i : F \rightarrow \mathbb{Z}$ such that the restriction of ϕ_i to $F \cap \langle \gamma_i \rangle$ is non-trivial but $\phi_i(f) = 0$ if $f \in F \cap \langle \delta \gamma_j \delta^{-1} \rangle$ for any $j \neq i$ and $\delta \in \Lambda$.*

Proof. We identify Λ with the fundamental group of a finite connected graph Y . By Marshall Hall's theorem, for each i there exists a finite-sheeted covering space $Y_i \rightarrow Y$ in which γ_i is represented by an embedded loop. Let $X \rightarrow Y$ be a regular, finite-sheeted covering space of Y that factors through Y_i for every i . Note that the generator of $\pi_1 X \cap \langle \delta \gamma_j \delta^{-1} \rangle$ is represented by an embedded loop in X , say $\lambda_{j,\delta}$, for all $j = 1, \dots, m$ and $\delta \in \Lambda$. In [41, Theorem 3.6], Wise proves that given any graph G and any simple loop λ in that graph, there is a finite-sheeted covering $\tilde{G} \rightarrow G$ in which any elevation of λ (i.e. a

lift of a power of λ) is independent in $H_1(\check{G}, \mathbb{Z})$ from the full set of elevations of all other simple loops in G . Applying this to the loop $\lambda_{i,1}$ in X , we obtain a finite-sheeted covering $X_i \rightarrow X$ and a homomorphism ψ_i from $F_i := \pi_1 X_i$ to \mathbb{Z} such that $\psi_i|_{F_i \cap \langle \gamma_i \rangle}$ is non-trivial but $\psi_i(F_i \cap \langle \delta \gamma_j \delta^{-1} \rangle) = 0$ for all $\delta \in \Lambda$ if $j \neq i$. Taking F to be the intersection of the F_i and $\phi_i = \psi_i|_F$ completes the proof. \square

We need to improve Lemma 4.4 to deal with virtually free groups Γ .

Lemma 4.5. *Let Γ be a virtually free group. If $(\gamma_1, \dots, \gamma_m)$ is an independent m -tuple, then there exists a free, normal subgroup $F < \Gamma$ of finite index and homomorphisms $\phi_i : F \rightarrow \mathbb{Z}$ such that the restriction of ϕ_i to $F \cap \langle \gamma_i \rangle$ is non-trivial but $\phi_i(f) = 0$ if $f \in F \cap \langle \delta \gamma_j \delta^{-1} \rangle$ for any $j \neq i$ and $\delta \in \Gamma$.*

Proof. By hypothesis, there is a short exact sequence of groups

$$1 \rightarrow \Lambda \rightarrow \Gamma \rightarrow \Sigma \rightarrow 1$$

with Λ free and Σ finite.

Given independent $\gamma_1, \dots, \gamma_m \in \Gamma$, we may replace the γ_i by proper powers to assume that each $\gamma_i \in \Lambda$. Then, we enlarge our list of elements by adding to it elements of Λ that are conjugate to some γ_i in Γ but not in Λ . To this end, we fix a set of coset representatives $\tilde{\Sigma} = \{\tilde{\sigma} \mid \sigma \in \Sigma\}$ for Λ in Γ , with $\tilde{1} = 1$, and define $g_{i\sigma} = \tilde{\sigma} \gamma_i \tilde{\sigma}^{-1}$. Since the γ_i are independent, no element of $\{g_{i\sigma} \mid \sigma \in \Sigma\}$ has a non-zero power that is conjugate to a non-zero power of an element of $\{g_{j\sigma} \mid \sigma \in \Sigma\}$ if $i \neq j$. However, the indexed set $(g_{i\sigma} \mid \sigma \in \Sigma)$ may fail to be independent since it is quite possible that $g_{i\sigma}$ will be conjugate to $g_{i\sigma'}^{\pm 1}$ for some $\sigma \neq \sigma'$. (In a virtually free group an element of infinite order x cannot be conjugate to x^p with $|p| > 1$, so higher powers are not a worry.) To account for such coincidences we make deletions from the list $(g_{i\sigma} \mid \sigma \in \Sigma)$, reducing it to $(g_{i\sigma} \mid \sigma \in \Sigma[i])$, say. This reduced list consists of a set of orbit representatives for the action of Γ by conjugation on the Λ -conjugacy classes of cyclic subgroups of the form $\langle f \gamma_i f^{-1} \rangle$ with $f \in \Gamma$.

We now apply Lemma 4.4 to the concatenation of the lists $(g_{1\sigma} \mid \sigma \in \Sigma[1]), \dots, (g_{m\sigma} \mid \sigma \in \Sigma[m])$, which is independent in Λ . Thus we obtain a free subgroup of finite index $F < \Lambda$ and homomorphisms $\phi_{i\sigma} : F \rightarrow \mathbb{Z}$ with the property that ϕ_{i1} is non-trivial on $F \cap \langle \gamma_i \rangle$ but $\phi_{i1}(f) = 0$ if $f \in F \cap \langle \delta g_{j\sigma} \delta^{-1} \rangle$ for any $\delta \in \Lambda$ and $(j, \sigma) \neq (i, 1)$. Moreover, since these conditions are inherited by subgroups of finite index in F , we may replace F by a smaller subgroup if necessary to ensure that it is normal in Γ .

Henceforth we write ϕ_i in place of ϕ_{i1} .

Consider $\delta \in \Gamma$ and γ_j with $j \neq i$. Write $\delta = \delta' \tilde{\sigma}$ for some $\delta' \in \Lambda$. If $\tilde{\sigma} \in \Sigma[j]$ then for any positive power n such that $\gamma_j^n \in F$ we have $\phi_i(\delta \gamma_j^n \delta^{-1}) = \phi_i(\delta' g_{j\sigma}^n (\delta')^{-1}) = 0$ as required. On the other hand, if $\tilde{\sigma} \notin \Sigma[j]$ then there exists $\lambda \in \Lambda$ such that $g_{j\sigma} = \lambda g_{j\sigma'}^{\pm 1} \lambda^{-1}$ for some $\sigma' \in \Sigma[j]$. Then,

$$\phi_i(\delta \gamma_j^n \delta^{-1}) = \phi_i(\delta' \lambda g_{j\sigma'}^{\pm n} \lambda^{-1} (\delta')^{-1}) = 0 ,$$

which finishes the proof. \square

With Lemma 4.5 in hand, we can prove Theorem 4.3.

Proof of Theorem 4.3. Let Γ , F and $\phi_i : F \rightarrow \mathbb{Z}$ be as in Lemma 4.5 and let $\eta : \Gamma \rightarrow \Gamma/F$ be the quotient map. Let $\ell_i = \phi_i(\gamma_i^{o(\eta(\gamma_i))})$ and note that there is no loss of generality in assuming that ℓ_i is positive. Fix a set of coset representatives c_j for F in Γ with $c_1 = 1$. For each γ_i , fix a positive integer N_i (to be specified later) and consider the composition

$$\psi_i : F \xrightarrow{\phi_i} \mathbb{Z} \rightarrow \mathbb{Z}/N_i .$$

Then, consider the direct product

$$\Psi_i = \prod_j \psi_i \circ i_{c_j} : F \rightarrow A_i = \prod_j \mathbb{Z}/N_i$$

where i_{c_j} is the automorphism of F given by conjugation by c_j . It is now clear that $o(\Psi_i(\gamma_i^{o(\eta(\gamma_i))})) = N_i/\ell_i$, whereas

$$\Psi_i(\gamma_k^{o(\eta(\gamma_k))}) = 0$$

for all $k \neq i$. The direct product

$$\Psi = \prod_i \Psi_i : F \rightarrow A = \prod_i A_i$$

therefore has the property that

$$o(\Psi(\gamma_i)^{o(\eta(\gamma_i))}) = N_i/\ell_i$$

for all i . Now, Ψ is the restriction to F of the homomorphism

$$\Phi : \Gamma \rightarrow A \rtimes (\Gamma/F) = \left(\prod_i \mathbb{Z}/N_i \right) \wr (\Gamma/F)$$

induced from $\prod_i \psi_i : F \rightarrow \prod_i \mathbb{Z}/N_i$. Therefore, $o(\Phi(\gamma_i)) = N_i o(\eta(\gamma_i))/\ell_i$ for all i .

To prove the theorem, we define Q to be $A \rtimes (\Gamma/F)$ and q to be Φ , then we take $\kappa = |\Gamma/F|^2$ and $N_i = \ell_i \kappa e_i / o(\eta(\gamma_i))$. The preceding computation shows that $o(\Phi(\gamma_i)) = \kappa e_i$, which proves the first assertion.

To prove the second assertion, suppose that an intersection,

$$\langle \Phi(\gamma_1) \rangle \cap \langle \Phi(\gamma_2) \rangle$$

say, is non-trivial. Then it contains a minimal non-trivial subgroup, which is of prime order p . That is, the intersection contains the non-trivial subgroup

$$\langle \Phi(\gamma_1^{\kappa e_1/p}) \rangle = \langle \Phi(\gamma_2^{\kappa e_2/p}) \rangle.$$

We have

$$o(\eta(\gamma_i)) \mid \kappa e_i/p$$

(because $\kappa = |\Gamma/F|^2$), and so $\gamma_i^{\kappa e_i/p} \in F$, for $i = 1, 2$. Therefore $\Psi(\gamma_i^{\kappa e_i/p}) = \Phi(\gamma_i^{\kappa e_i/p})$ for $i = 1, 2$, and so

$$\langle \Psi(\gamma_1^{\kappa e_1/p}) \rangle = \langle \Psi(\gamma_2^{\kappa e_2/p}) \rangle.$$

One of the coordinates of the homomorphism Ψ is ψ_1 , and so it follows that

$$\langle \psi_1(\gamma_1^{\kappa e_1/p}) \rangle = \langle \psi_1(\gamma_2^{\kappa e_2/p}) \rangle.$$

But this leads to a contradiction because, on the one hand, we have $\psi_1(\gamma_2^{\kappa e_2/p}) = 0$ by the definition of ψ_1 and Lemma 4.5, while on the other hand, $\psi_1(\gamma_1^{\kappa e_1/p}) \neq 0$, because

$$\psi_1(\gamma_1^{\kappa e_1/p}) = \psi_1(\gamma_1^{o(\eta(\gamma_1)) \kappa e_1/p o(\eta(\gamma_1))})$$

and $\kappa e_1/p o(\eta(\gamma_1))$ is less than $o(\psi_1(\gamma_1^{o(\eta(\gamma_1))})) = (N_1/\ell_1) = \kappa e_1/o(\eta(\gamma_1))$. \square

5. CONSTRUCTING MALNORMAL SUBGROUPS

The role that malnormality plays in our strategy was explained in Section 3. The main result in this section is Proposition 5.9, but several of the other lemmas will also be required in the next section. Fibre products of morphisms of graphs, as described by Stallings [37], play a prominent role in many of our proofs.

Definition 5.1. Let Γ be a group and H a subgroup. Then H is said to be *almost malnormal* in Γ if $|H \cap H^\gamma| < \infty$ whenever $\gamma \in \Gamma \setminus H$. If we in fact have $H \cap H^\gamma = 1$ whenever $\gamma \in \Gamma \setminus H$ then H is said to be *malnormal*.

More generally, a family $\{H_i\}$ of subgroups of Γ is said to be *almost malnormal* if $|H_i \cap H_j^\gamma| = \infty$ implies that $i = j$ and $\gamma \in H_j$. Similarly, we may speak of malnormal families of subgroups.

Note that if H is torsion-free and almost malnormal then it is in fact malnormal.

The first fact we record is trivial but extremely useful.

Lemma 5.2. *If K is an (almost) malnormal subgroup of H and H is an almost malnormal subgroup of G then K is an (almost) malnormal subgroup of G .*

The next lemma, which again admits a trivial proof, enables one to deduce almost malnormality from virtual considerations.

Lemma 5.3. *Let H be an arbitrary subgroup of a group Γ and let Γ_0 be a subgroup of finite index in Γ . Fix a set of double-coset representatives $\{\gamma_i\}$ for $H\backslash\Gamma/\Gamma_0$. Then H is almost malnormal in Γ if and only if the family $\{H^{\gamma_i} \cap \Gamma_0\}$ is almost malnormal in Γ_0 .*

The malnormality of a family of subgroups of a free group can be determined by a computation using the elegant formalism of fibre products, as we will now explain.

Consider a pair of immersions of finite graphs $\iota_1 : Y_1 \rightarrow X$ and $\iota_2 : Y_2 \rightarrow X$. Recall that the *fibre product* of the maps ι_1 and ι_2 is defined to be the graph

$$Y_1 \times_X Y_2 = \{(y_1, y_2) \in Y_1 \times Y_2 \mid \iota_1(y_1) = \iota_2(y_2)\}.$$

The fibre product comes equipped with a natural immersion $\kappa : Y_1 \times_X Y_2 \rightarrow X$. For any (y_1, y_2) , Stallings pointed out that

$$\kappa_*\pi_1(Y_1 \times_X Y_2, (y_1, y_2)) = \iota_{1*}\pi_1(Y_1, y_1) \cap \iota_{2*}\pi_1(Y_2, y_2)$$

[37, Theorem 5.5]. In the case when $Y_1 = Y_2 = Y$ and $\iota_1 = \iota_2$, there is a canonical diagonal component of $Y \times_X Y$, isometric to Y .

The next lemma follows immediately from this discussion.

Lemma 5.4. *Let X be a connected finite graph with fundamental group F , and let Y be a (not necessarily connected) finite graph equipped with an immersion $Y \rightarrow X$. The components $\{Y_i\}$ of Y define (up to conjugacy) a family of subgroups H_i of F . Then $\{H_i\}$ is malnormal if and only if every non-diagonal component of the fibre product $Y \times_X Y$ is simply connected.*

In particular, this gives an algorithm to determine whether or not a given family of subgroups of a free group is malnormal.

Unlike Lemma 5.4, the next lemma is not always applicable. However, it gives a useful sufficient condition for malnormality, which can sometimes be applied in situations where Lemma 5.4 is too cumbersome to apply in practice. Let $Z_\Gamma(g)$ denote the centralizer of an element g in a group Γ .

Lemma 5.5. *Let H be a subgroup of Γ . If H is a retract and $Z_\Gamma(h) \subseteq H$ for all $h \in H \setminus 1$, then H is malnormal.*

Proof. Let $\rho : \Gamma \rightarrow H$ be a retraction. Suppose that $h \in H \setminus 1$ and $h^\gamma \in H$. Then $h^\gamma = h^{\rho(\gamma)}$, which implies that $\gamma\rho(\gamma)^{-1} \in Z_\Gamma(h)$ and so $\gamma \in H$, as required. \square

We now develop some simple examples.

Example 5.6. If Γ is a group and H is a free factor then H is malnormal in Γ . This is an immediate consequence of Lemma 5.5, since free factors are retracts.

The following easy example, which will be useful later, illustrates how Lemmas 5.3 and 5.4 can be used to prove almost malnormality in virtually free groups.

Example 5.7. Suppose that A is a finite group and B is any subgroup of A . Then the natural copy of $H = B * \mathbb{Z}$ inside $\Gamma = A * \mathbb{Z}$ is almost malnormal.

To see this, realize Γ as the fundamental group of a graph of groups \mathcal{X} with a single vertex labelled A and a single edge with trivial edge group. The kernel of the retraction $\Gamma \rightarrow A$ implicit in the notation is a normal, free subgroup F of finite index. Let T be the Bass–Serre tree of \mathcal{X} . The quotient $F \backslash T$ is a graph X with a single vertex and $|A|$ edges $\{e_a \mid a \in A\}$, and the natural A -action is by left translation. The subgroup $H \cap F$ is carried by the subgraph $Y = \bigcup_{b \in B} e_b$.

The quotient map $\Gamma \rightarrow A$ identifies $H \backslash \Gamma / F$ with $B \backslash A$, so a set of double-coset representatives for the former is provided by any set $\{a_i\}$ of right-coset representatives for B in A . The subgroup $H^{a_i} \cap F$ is carried by the subgraph $a_i^{-1}Y$: under the immersion

$$Z = \coprod_i a_i^{-1}Y \rightarrow X$$

(where the map $Z \rightarrow X$ is inclusion on each component), the fundamental groups of the components are mapped to the family of subgroups $\{H^{a_i} \cap F\}$.

Note that, as subgraphs of X , $a_i^{-1}Y$ and $a_j^{-1}Y$ have no edges in common if $i \neq j$. Therefore, every off-diagonal component of $Z \times_X Z$ is a vertex and hence simply connected.

It follows that $\{H^{a_i} \cap F\}$ forms a malnormal family in F by Lemma 5.4, and so H is almost malnormal in Γ by Lemma 5.3.

The following construction provides us with the supply of malnormal subgroups that we shall need to prove Theorem A.

Lemma 5.8. *Let $\Lambda_2 \cong \langle \alpha, \beta \rangle$ be free of rank two. For each integer N , let*

$$q_N : \Lambda_2 \rightarrow Q_N = \Lambda_2 / \langle\langle \beta^N \rangle\rangle$$

be the quotient map. Consider $u = \alpha^\beta(\alpha^{\beta^2})^{-1}$, $v = \alpha^\beta(\alpha^{\beta^{-1}})^{-1}$. For all $N > 6$, the subgroup $q_N(\langle \alpha, u, v \rangle)$ is malnormal in Q_N and free of rank 3.

Proof. Consider the images $\bar{\alpha} = q_N(\alpha)$, $\bar{\beta} = q_N(\beta)$, $\bar{u} = q_N(u)$ and $\bar{v} = q_N(v)$. Let F be the kernel of the retraction $Q_N \rightarrow \mathbb{Z}/N$ that maps $\bar{\alpha} \mapsto 0$ and $\bar{\beta} \mapsto 1$. As above, F may be thought of as the fundamental group of a graph X with a single vertex, and with N edges $\{e_i\}_{i \in \mathbb{Z}/N}$, on which $\mathbb{Z}/N = \langle \bar{\beta} \rangle$ acts by left translation. Represent $\langle \bar{\alpha}, \bar{u}, \bar{v} \rangle$ by the usual immersion of core graphs $\iota : Y \rightarrow X$. As long as $N \geq 4$, the core graph Y is easily computed explicitly using Stallings folds (see Figure 1), and is seen to have rank 3 as required.




FIGURE 1. The core graph Y . The edges are labelled with their images in X .

Let $Y_i = Y$ for each $i = 0, \dots, N-1$, and consider the disjoint union

$$Z = \coprod_{i=0}^{N-1} Y_i \rightarrow X$$

where the map on Y_i is $\bar{\beta}^i \circ \iota$. To prove malnormality, it suffices to argue that every off-diagonal component of the fibre product $Z \times_X Z$ is simply connected.

Suppose some off-diagonal component is not simply connected. Translating by an element of $\langle \bar{\beta} \rangle$, we may assume that it arises as part of the fibre product $Y_0 \times_X Y_i$ for some i . Since the image of ι only contains the edges e_i for $-1 \leq i \leq 2$, this fibre product contains no edges unless $0 \leq i \leq 3$ (because $N > 6$).

Therefore, it is enough to check that the off-diagonal components of $Y_0 \times_X Y_0$ are simply connected, and that every component of $Y_0 \times_X Y_i$ is simply connected, where $i = 1, 2, 3$. The off-diagonal components of $Y_0 \times_X Y_0$ are points; for $i = 1, 2, 3$, the fibre product $Y_0 \times_X Y_i$ has $4 - i$ edges, and a direct computation shows that each of these is a forest. The fibre products $Y_0 \times_X Y_0$ and $Y_0 \times_X Y_1$ are illustrated in Figure

2, while $Y_0 \times_X Y_2$ and $Y_0 \times_X Y_3$ are left as easy computations for the reader. \square




FIGURE 2. The fibre products $Y_0 \times_X Y_0$ and $Y_0 \times_X Y_1$, displayed as subsets of the direct products $Y_0 \times Y_0$ and $Y_0 \times Y_1$. Note that the only non-simply-connected component is the diagonal component of $Y_0 \times_X Y_0$.

From the 3-generator case, we immediately obtain malnormal subgroups with arbitrarily many generators.

Proposition 5.9. *Let $\Lambda_2 \cong \langle \alpha, \beta \rangle$ be free of rank two. For each integer N , let*

$$q_N : \Lambda_2 \rightarrow Q_N = \Lambda_2 / \langle\langle \beta^N \rangle\rangle$$

be the quotient map. For any m , there exist $\{\gamma_0, \dots, \gamma_{m+1}\} \in [\Lambda_2, \Lambda_2]$ such that, for all $N > 6$, the subgroup $q_N(\langle \alpha, \gamma_0, \dots, \gamma_{m+1} \rangle)$ is malnormal in Q_N and free of rank $m + 3$.

Proof. Let $L = \langle \gamma_0, \dots, \gamma_{m+1} \rangle$ be any rank- $(m+2)$ malnormal subgroup of the free group $\langle u, v \rangle$ constructed in Lemma 5.8. Then $L * \langle \alpha \rangle$ is malnormal in $\langle \alpha, u, v \rangle$ and hence in Q_N , by Lemma 5.2. Also, since u and v lie in $[\Lambda_2, \Lambda_2]$, the γ_i do as well. \square

6. THE PROOF OF THEOREM C

In this section we prove Theorem C, following the strategy laid out in Section 3. As mentioned in the introduction, Theorem A follows immediately, using Theorem 2.1.

We are given a finitely presented group $G = \langle A \mid R \rangle = \langle a_1, \dots, a_m \mid r_1, \dots, r_n \rangle$ and a word $w \in F(A)$.

Step 1: improving the input. We start by proving some lemmas that improve the input G and w .

Lemma 6.1. *There is an algorithm that takes as input a finitely presented group $G \cong \langle A \mid R \rangle$ and a word $w \in F(A)$ and outputs a finite presentation $\langle A^\dagger \mid R^\dagger \rangle$ for a group G^\dagger and a word $w^\dagger \in F(A^\dagger)$ such that:*

- (1) $w^\dagger =_{\widehat{G}^\dagger} 1$ if and only if $w =_{\widehat{G}} 1$;
- (2) if $w^\dagger \neq_{\widehat{G}^\dagger} 1$ then the natural map $\{1\} \sqcup A^\dagger \rightarrow \widehat{G}^\dagger$ is an embedding.

Proof. Take $2m + 1$ copies $G^{(j)}$ of G and let a_{ij} be the copy of a_i in $G^{(j)}$; similarly, let w_j be the copy of w in $G^{(j)}$. We will always take the j index modulo $2m + 1$. Set

$$G^\dagger = G^{(1)} * \dots * G^{(2m+1)}$$

and note that $w_j =_{\widehat{G}^\dagger} 1$ if and only if $w =_{\widehat{G}} 1$. Now consider the following generating set A^\dagger for G^\dagger :

$$\{a_{ij}w_{j+m+1}w_{i+j} \mid 1 \leq i \leq m, 1 \leq j \leq 2m+1\} \cup \{w_j \mid 1 \leq j \leq 2m+1\}.$$

Let $G \rightarrow \Sigma$ be a finite quotient in which w survives, let $\eta_j : G^{(j)} \rightarrow \Sigma^{(j)}$ be the corresponding quotient of $G^{(j)}$ and consider the free product of the maps η_j

$$\eta^\dagger : G^\dagger \rightarrow \Sigma^\dagger = \Sigma^{(1)} * \dots * \Sigma^{(2m+1)}.$$

Suppose now that

$$\eta^\dagger(a_{ij}w_{j+m+1}w_{i+j}) = \eta^\dagger(a_{i'j'}w_{j'+m+1}w_{i'+j'})$$

for some i, j, i', j' . Because $i, i' < m + 1$, the three syllables of the product $a_{ij}w_{j+m+1}w_{i+j}$ lie in different free factors, hence $j = j'$ and $i = i'$. Similarly, the images of the generators w_j lie in unique and distinct free factors. Therefore, the restriction of η^\dagger to $1 \sqcup A^\dagger$ is injective. Since Σ^\dagger is virtually free and hence residually finite, it follows that $1 \sqcup A^\dagger$ injects into \widehat{G}^\dagger as required. Setting $w^\dagger = w_1$ finishes the proof. \square

Proposition 6.2. *There is an algorithm that takes as input a finitely presented group $G \cong \langle A \mid R \rangle$ and a word $w \in F(A)$ and outputs a finite presentation $\langle A' \mid R' \rangle$ for a group G' and a word $w' \in F(A')$ such that:*

- (1) $w' =_{\widehat{G}'} 1$ if and only if $w =_{\widehat{G}} 1$;
- (2) if $w \neq_{\widehat{G}} 1$ then, for any $N \in \mathbb{N}$, there exists a homomorphism to a finite group $\eta : G' \rightarrow Q$ such that:
 - (a) $o(\eta(a')) = o(\eta(w')) \geq N$ for all $a' \in A'$; and

(b) $\langle \eta(a'_i) \rangle \cap \langle \eta(a'_j) \rangle = \langle \eta(a'_i) \rangle \cap \langle \eta(w') \rangle = 1$ whenever $i \neq j$.

Proof. We may algorithmically construct a presentation $\langle A^\dagger \mid R^\dagger \rangle$ and a word w^\dagger as in Lemma 6.1. Write $A^\dagger = \{a_1^\dagger, \dots, a_m^\dagger\}$. Let $G' = G^\dagger * \langle a'_0 \rangle$, let $a'_i = a_i^\dagger a'_0$ for each i and let $w' = [w^\dagger, a'_0]$. Assertion (1) is now immediate.

Let $\eta^\dagger : G^\dagger \rightarrow \Sigma^\dagger$ be as in the proof of Lemma 6.1 and let $\Gamma = \Sigma^\dagger * \langle a_0 \rangle$. We extend η to a surjection $\zeta : G' \rightarrow \Gamma$ by defining $\zeta(a_0) = a_0$. The map η is injective on $1 \sqcup A^\dagger$, so by the normal form theorem for free products, $(\zeta(a'_0), \dots, \zeta(a'_m), \zeta(w'))$ is an independent $(m+2)$ -tuple in Γ . To complete the proof, we define η to be the composition of ζ and the map $q : \Gamma \rightarrow Q$ provided by Theorem 4.3. \square

To avoid being overwhelmed by notation, we rename G' as G , $A' = \{a'_0, \dots, a'_m\}$ as A and w' as w .

Step 2: a map $G \rightarrow \hat{G}_1$ whose image is trivial iff $w = 1$. We define a new finitely presented group

$$G_1 = G * \langle b_0, \dots, b_m \rangle / \langle\langle w^{b_i} = a_i \mid i = 0, \dots, m \rangle\rangle$$

and let F_0 denote the subgroup $\langle b_0, \dots, b_m \rangle$. Note that there is a retraction $\rho : G_1 \rightarrow F_0$, whence F_0 is free of rank $m+1$. Note too that there is a simple algorithm for deriving a finite presentation of G_1 from G and w . The following lemma is clear.

Lemma 6.3. *If $w =_{\hat{G}} 1$, then the inclusion map $F_0 \hookrightarrow G_1$ and the retraction ρ induce isomorphisms of profinite completions.*

If $w \neq_{\hat{G}} 1$ then we have the finite quotient $\eta : G \rightarrow Q$ guaranteed by Proposition 6.2. We extend η to an epimorphism from G_1 to the virtually free group Γ_0 given by the relative presentation below. We will continue to denote this epimorphism by η and, to further simplify notation, we will use bars to denote the image of an element or a subgroup under η , so $\eta(w) = \bar{w}$, $\eta(F_0) = \bar{F}_0$ etc.

$$\Gamma_0 = (Q, \bar{b}_1, \dots, \bar{b}_m \mid \bar{w}^{\bar{b}_i} = \bar{a}_i \text{ for } i = 0, \dots, m)$$

Note that $o(\bar{w}) = o(\bar{a}_i)$ in Q , by Proposition 6.2, and therefore Γ_0 is a multiple HNN extension of Q . Let \mathcal{X}_0 be the corresponding graph of groups and let T_0 be its Bass–Serre tree.

Lemma 6.4. *If $w \neq_{\hat{G}} 1$ then, for all natural numbers N , the group G_1 has a virtually free quotient $\eta : G_1 \rightarrow \Gamma_0$ with the following properties:*

- (1) *for all $a \in A$, $N \leq o(\bar{w}) = o(\bar{a}) < \infty$;*
- (2) *\bar{F}_0 is free of rank $m+1$ and malnormal in Γ_0 .*

Proof. The map $\eta : G_1 \rightarrow \Gamma_0$ was constructed above. In the light of Proposition 6.2, the only point that is not immediate is that \overline{F}_0 is malnormal. The quotient of Γ_0 by Q defines a retraction $\bar{\rho}$ from Γ_0 to \overline{F}_0 . By Lemma 5.5, it suffices to prove that $Z_{\Gamma_0}(h) \subseteq \overline{F}_0$ for all $h \in \overline{F}_0 \setminus 1$.

Suppose therefore that $h \in \overline{F}_0 \setminus 1$ and $[h, \gamma] = 1$. Let $h = \bar{b}_{i_1}^{\epsilon_1} \dots \bar{b}_{i_k}^{\epsilon_k}$, where $\epsilon_i \in \{\pm 1\}$ for all i . We may assume that this decomposition is cyclically reduced, and therefore the vertex $*$ in T_0 stabilized by Q is on the axis of h .

We claim that $\bar{\rho}(\gamma^{-1})\gamma \in Q$. Because $[h, \gamma] = 1$, the segment $[*, \gamma*]$ is contained in $\text{Axis}(h)$. Because \overline{F}_0 acts on its minimal invariant subtree with a single orbit of vertices, there exists $\beta \in \overline{F}_0$ such that $\beta* = \gamma*$, and so $\beta^{-1}\gamma \in Q$. Therefore

$$1 = \bar{\rho}(\beta^{-1}\gamma) = \beta^{-1}\bar{\rho}(\gamma)$$

and the claim follows.

Since $\bar{\rho}(\gamma) \in Z_{\Gamma_0}(h)$, the claim reduces us to the case that $\gamma \in Q$, in which case γ fixes the whole of $\text{Axis}(h)$. But, by item (2)(b) of Proposition 6.2, no non-trivial element of Γ_0 fixes a subset of diameter greater than 2 in the minimal \overline{F}_0 -invariant subtree of T_0 , and therefore $\gamma = 1$. \square

Step 3: the free subgroups F_1 , F_2 and F . Let $G_2 = G_1 * \langle t \rangle$ and let F_1 be the free subgroup $F_0 * \langle t \rangle$ of rank $m + 2$. It will later be convenient to write $b_{m+1} = t$. Casting t and w in the roles of α and β , we choose $c_j = \gamma_j$ as in Proposition 5.9, for $j = 0, \dots, m + 1$, and write F_2 for the subgroup of G_2 generated by the c_j .

Since c_j is in the commutator subgroup of $\langle t, w \rangle$, we have $c_j \in \langle\langle w \rangle\rangle$ and hence the image of F_2 in \widehat{G}_2 is trivial if $w =_{\widehat{G}} 1$.

We analyse what happens when $w \neq_{\widehat{G}} 1$. Let $\eta : G_1 \rightarrow \Gamma_0$ be the virtually free quotient guaranteed by Lemma 6.4. We will extend η to a homomorphism from G_2 onto a virtually free group Γ ; we will then continue to denote this homomorphism by η , and continue to denote η -images by bars.

Consider the graph of groups \mathcal{X} obtained from \mathcal{X}_0 by adjoining a single loop e_{m+1} with trivial edge group; denote the corresponding stable letter by \bar{t} (it will also sometimes be convenient to denote it by \bar{b}_{m+1}). We define

$$\Gamma = \pi_1 \mathcal{X}$$

and extend η to $\eta : G_2 \rightarrow \Gamma$ by setting $\eta(t) = \bar{t}$.

Let $F = \langle F_1, F_2 \rangle$. The remainder of this section is devoted to an analysis of the image $\eta(F) = \overline{F} \subseteq \Gamma$.

Let $J_0 \triangleleft \Gamma_0$ be a normal, free subgroup of finite index. We use the canonical retraction $\bar{\rho} : \Gamma_0 \rightarrow \bar{F}_0$ to modify J_0 , replacing it with $K_0 = J_0 \cap \bar{\rho}^{-1}(J_0 \cap \bar{F}_0)$. The quotient $K_0 \backslash T_0$ is a graph X_0 with fundamental group K_0 ; X_0 may be thought of as a finite-sheeted covering space of the graph of groups \mathcal{X}_0 (this can be made formal, but we will avoid using it explicitly). There is a natural vertex-transitive left-action of $P = \Gamma_0/K_0$ on X_0 , in which the stabilizer of each vertex is conjugate to Q (note that Q embeds into P since $Q \cap K_0 = 1$). In particular, fixing a base vertex $*$ for X_0 , we may identify the vertex set of X_0 with the coset space P/Q .

There is a minimal \bar{F}_0 -invariant subtree $T_0^{\bar{F}_0} \subseteq T_0$. Let $Y_0 = (\bar{F}_0 \cap K_0) \backslash T_0^{\bar{F}_0}$. The inclusion map descends to a combinatorial map $Y_0 \rightarrow X_0$. Picking a base vertex in Y_0 , this map represents the inclusion $\bar{F}_0 \cap K_0 \rightarrow \Gamma_0$. In fact, this map is an embedding.

Lemma 6.5. *The graph Y_0 is a regular covering of the rose with $m+1$ petals, and the map $\iota : Y_0 \rightarrow X_0$ is an embedding.*

Proof. Note that \bar{F}_0 acts freely on $T_0^{\bar{F}_0}$ and transitively on the vertices. Therefore, the quotient $\bar{F}_0 \backslash T_0^{\bar{F}_0}$ is the rose with $m+1$ petals, and $(\bar{F}_0 \cap K_0) \backslash T_0^{\bar{F}_0}$ is a regular covering with deck group $R := \bar{F}_0/(\bar{F}_0 \cap K_0)$.

By standard Bass–Serre theory, the fact that $Y_0 \rightarrow X_0$ is an embedding reduces to the fact that the natural maps

$$(\bar{F}_0 \cap K_0) \backslash \bar{F}_0 \rightarrow K_0 \backslash \Gamma_0 / Q$$

and

$$(\bar{F}_0 \cap K_0) \backslash \bar{F}_0 \rightarrow K_0 \backslash \Gamma_0 / \langle \bar{w} \rangle$$

are injective. (More exactly, the injectivity of the first map above implies the injectivity of $Y_0 \rightarrow X_0$ on vertices, and the injectivity of the second map implies the injectivity of $Y_0 \rightarrow X_0$ on edges.) Since the first map factors through the second, it is enough to prove that $(\bar{F}_0 \cap K_0) \backslash \bar{F}_0 \rightarrow K_0 \backslash \Gamma_0 / Q$ is injective.

Suppose therefore that $f, g \in \bar{F}_0$ and $f = kgq$, where $k \in K_0$ and $q \in Q$. Applying the retraction $\bar{\rho} : \Gamma_0 \rightarrow \bar{F}_0$, we deduce that $f = \bar{\rho}(k)g$, which implies that

$$q = (k^{-1}\bar{\rho}(k))^g.$$

But q has finite order and $k^{-1}\bar{\rho}(k)$ lies in the free group J_0 . Therefore $q = 1$ and

$$(\bar{F}_0 \cap K_0)f = (\bar{F}_0 \cap K_0)g$$

as required. \square

We will identify Y_0 with its image in X_0 , and hence we feel free to (without loss of generality) choose $*$ as the base point for Y_0 . Fixing a base point allows us to identify the vertices of Y_0 with the elements of R .

There is a natural retraction $\sigma : \Gamma \rightarrow \Gamma_0$ obtained by setting $\sigma(\bar{t}) = 1$. The preimage $K = \sigma^{-1}(K_0)$ is a normal, free subgroup of finite index in Γ with $\Gamma/K \cong P$. Let T be the Bass–Serre tree of \mathcal{X} . Then $X = K \backslash T$ is a finite graph which, as before, can be thought of as a regular, finite-sheeted covering space of \mathcal{X} with deck group P .

In fact, there is a simple, concrete description of X . Consider the graph of groups \mathcal{Z} with a single vertex, labelled by the finite group Q , and a single edge, with trivial edge group. Its fundamental group is $Q * \mathbb{Z}$, which can be identified with $Q * \langle \bar{t} \rangle$, a subgroup of Γ . There is an obvious retraction $Q * \langle \bar{t} \rangle \rightarrow Q$ obtained by sending $\bar{t} \mapsto 1$, and the kernel is precisely $(Q * \langle \bar{t} \rangle) \cap K$, a normal, torsion-free subgroup of finite index, with quotient group Q . The corresponding covering graph of \mathcal{Z} can be constructed as follows. Let Z be the graph with one vertex and edges $\{e_q \mid q \in Q\}$. This admits a natural Q -action, where Q acts freely on the edges e_q by left translation, and its fundamental group can be identified with $(Q * \langle \bar{t} \rangle) \cap K$.

For each coset $pQ \in P/Q$, let Z^{pQ} be a copy of Z . Now X can be constructed as a quotient

$$X = \left(X_0 \sqcup \coprod_{pQ \in P/Q} Z^{pQ} \right) / \sim$$

where \sim identifies the unique vertex of Z^{pQ} with the vertex of X_0 that corresponds to pQ (i.e. $p*$). The group P acts on X ; the vertex $p*$ is stabilized by $Q^{p^{-1}}$, which acts freely on the edges of Z^{pQ} .

The inclusion $Y_0 \rightarrow X_0$ provides us with a nice geometric representative for the inclusion of $\bar{F}_0 \cap K_0$ into K_0 . We next extend this to a nice geometric representative for $\bar{F} \cap K$ in K .

Let $W \rightarrow Z$ be an immersion (with basepoints) representing $\langle \bar{t} \rangle * \bar{F}_2 = \langle \bar{t}, \bar{c}_0, \dots, \bar{c}_{m+1} \rangle$ as a subgroup of the kernel of the natural retraction $Q * \langle \bar{t} \rangle \rightarrow Q$. (Note that this immersion exists because $\bar{F}_2 \subseteq \langle\langle \bar{t} \rangle\rangle$.) Take copies $W^p \equiv W$, one for each $p \in P$, equipped with maps $W^p \rightarrow Z^{pQ}$, chosen so that if $pQ = p'Q$ then the following diagram commutes:

$$\begin{array}{ccc} W^p & \longrightarrow & Z^{pQ} \\ \downarrow \equiv & & \downarrow p'p^{-1} \\ W^{p'} & \longrightarrow & Z^{pQ} \end{array}$$

where we note that $p'p^{-1} = (p^{-1}p')^{p^{-1}} \in Q^{p^{-1}}$, which acts on Z^{pQ} as remarked above. Now let

$$Y = \left(Y_0 \sqcup \coprod_{r \in R} W^r \right) / \sim$$

where \sim identifies the base vertex of W^r with the vertex $r* \in Y_0$. (Recall that $R = \overline{F}_0 / (\overline{F}_0 \cap K_0)$, i.e. the image of \overline{F}_0 in P , as in the proof of Lemma 6.5.) The coproduct of the embedding $Y_0 \hookrightarrow X_0$ and the immersions $W^r \rightarrow Z^{pQ}$ is an immersion $Y \rightarrow X$, since adjacent edges of Y_0 and W^r map to distinct edges of X . Taking $* \in Y_0$ as a base vertex for Y , the immersion $Y \rightarrow X$ represents the inclusion of $\overline{F} \cap K$ into K .

Lemma 6.6. *If $w \neq_{\widehat{G}} 1$ then $\overline{F} = \overline{F}_1 * \overline{F}_2$.*

Proof. Because free groups are Hopfian, it suffices to prove that $\text{rk } \overline{F} = \text{rk } \overline{F}_1 + \text{rk } \overline{F}_2$. This can be deduced from a computation of the Euler characteristic of Y , as follows.

If $d = |R|$, then we have

$$\begin{aligned} \chi(Y) &= \chi(Y_0) + d\chi(W) - d \\ &= d((1 - \text{rk } \overline{F}_0) + (1 - (1 + \text{rk } \overline{F}_2)) - 1) \\ &= d(1 - (\text{rk } \overline{F}_0 + 1 + \text{rk } \overline{F}_2)) \\ &= d(1 - (2m + 4)) . \end{aligned}$$

On the other hand, the fundamental group of Y is $K \cap \overline{F}$, which is of index d in \overline{F} . Therefore

$$\chi(Y) = d(1 - \text{rk } \overline{F}) .$$

So $\text{rk } \overline{F} = 2m + 4$, which is equal to $\text{rk } \overline{F}_1 + \text{rk } \overline{F}_2$. \square

Malnormality of \overline{F} . We shall establish the malnormality of \overline{F} using the immersion $Y \rightarrow X$. For each left coset $pR \in P/R$, let Y_0^{pR} be a copy of Y_0 . For each coset pR we choose a representative p_i and equip Y_0^{pR} with the inclusion in X_0 that is the composition of p_i with the inclusion $Y_0 \rightarrow X_0$.

Consider

$$U_0 = \coprod_{pR \in P/R} Y_0^{pR} \rightarrow X_0 ,$$

the coproduct of the maps described above. There is a free action of the group P on U_0 obtained by insisting that R acts on Y_0^R in the usual way and that p_i takes the base vertex $*_R \in Y_0^R$ to the base vertex $*_{p_i R} \in Y_0^{p_i R}$, and with this definition the map $U_0 \rightarrow X_0$ is P -equivariant. Thus, the

vertices of U_0 are in bijection with the elements of P . The vertices of X_0 are in bijection with P/Q , and under this correspondence the map $U_0 \rightarrow X_0$ on the vertices can be seen as the natural map $P \rightarrow P/Q$.

Remark 6.7. Consider the fibre product $U_0 \times_{X_0} U_0$. Note that the map $U_0 \rightarrow X_0$ represents the family of subgroups $\{\overline{F}_0^{\gamma_i^{-1}} \cap K_0\}$ in K_0 , where γ_i ranges over a set of representatives for $K_0 \backslash \Gamma_0 / \overline{F}_0$ (which is identified with P/R). Therefore, by Lemmas 5.3, 5.4 and 6.4, the off-diagonal components of $U_0 \times_{X_0} U_0$ are simply connected.

We now consider the same construction for $Y \rightarrow X$. Let $Y^{pR} = Y$ and consider the disjoint union

$$U = \coprod_{pR \in P/R} Y^{pR} \rightarrow X$$

where, as before, the map $Y^{pR} \rightarrow X$ is the composition of a choice of map $p : X \rightarrow X$ with the immersion $Y \rightarrow X$. Alternatively, we can construct U from U_0 by attaching copies of W as follows:

$$U = \left(U_0 \sqcup \coprod_{p \in P} W^p \right) / \sim$$

where \sim identifies the vertex $p *_R \in U_0$ with the base vertex of W^p .

The map $U \rightarrow X$ represents the family of subgroups $\{\overline{F}^{\gamma_i^{-1}} \cap K\}$ in K , where γ_i ranges over a set of representatives for $K \backslash \Gamma / \overline{F} = P/R$; therefore, we will be able to prove the malnormality of \overline{F} by considering the fibre product $U \times_X U$.

We can obtain a clearer picture of the map $U \rightarrow X$ by first gathering together those copies of W whose images adjoin the same vertex of X . Let

$$V = \bigcup_{q \in Q} W^q \subseteq U$$

and note that

$$U = U_0 \cup \bigcup_{p_i Q \in P/Q} p_i V .$$

Then $p_i V$ is precisely the preimage of $Z^{p_i Q} \subseteq X$ under the map $U \rightarrow X$.

Lemma 6.8. *If $N > 6$ then the off-diagonal components of $V \times_Z V$ are simply connected.*

Proof. By Lemmas 5.3 and 5.4, this is equivalent to the claim that $\langle \bar{t} \rangle * \overline{F}_2 \subseteq \langle \bar{t} \rangle * \langle \bar{w} \rangle$ is malnormal in $\langle \bar{t} \rangle * Q$. This follows from Proposition 5.9, Example 5.7 and Lemma 5.2. \square

The fibre product $U \times_X U$ decomposes as

$$U \times_X U = (U_0 \times_{X_0} U_0) \cup \coprod_{p_i Q \in P/Q} (p_i V \times_{Z^{p_i Q}} p_i V)$$

and the diagonal components of $U \times_X U$ consist of precisely the diagonal components of the fibre products on the right hand side of the equation.

Proposition 6.9. *If $N > 6$ and $w \neq_{\hat{G}} 1$ then \overline{F} is malnormal in Γ .*

Proof. By Lemmas 5.3 and 5.4, it suffices to show that every off-diagonal component of the fibre product $U \times_X U$ is simply connected.

Suppose therefore that δ is a geodesic loop in an off-diagonal component of U . The fibre product is equipped with two projections $\pi_1, \pi_2 : U \times_X U \rightarrow U$ and a P -action. Let $\delta_i = \pi_i \circ \delta$. Translating by an element of P , we may assume that δ_1 is contained in Y^R .

If $\delta_1 \subseteq Y_0^R \subseteq Y^R$ then $\delta_2 \subseteq Y_0^{pR} \subseteq Y^{pR}$ for some $p \in P$, so δ is an essential off-diagonal loop in $U_0 \times_{X_0} U_0$, which contradicts the fact that \overline{F}_0 is malnormal in Γ . Therefore, δ_1 has a non-trivial subpath contained in W^r for some $r \in R$. Let α_1 be a maximal such subpath, let α be the subpath of δ with $\pi_1 \circ \alpha = \alpha_1$ and let $\alpha_2 = \pi_2 \circ \alpha$.

The endpoints of α_1 lie in $W^r \cap Y_0^R \subseteq Y$; this intersection is a point, and hence α_1 is a loop in W^r . Likewise, the endpoints of α_2 lie in $W^p \cap Y_0^{pR}$, which is also a point, and so α_2 is a loop in W^p . Since they have the same image in X it follows that $p = rq$ for some $q \in Q$. The loop $r^{-1}\delta$ is then a non-trivial loop in an off-diagonal component of $V \times_Z V$, which contradicts Lemma 6.8 (since $N > 6$). \square

Step 4: the end of the proof of Theorem C. We take two copies of G_2 , distinguishing elements and subgroups of the second by primes, and define G_w to be the quotient of $G_2 * G'_2$ by the relations

$$\{c_i = b'_i, b_i = c'_i \mid i = 0, \dots, m+1\}.$$

If $w =_{\hat{G}} 1$, it is clear that $\hat{G}_w \cong 1$.

Suppose that $w \neq_{\hat{G}} 1$. Then G_w is the amalgamated product

$$G_2 *_F \cong_{F'} G'_2$$

where the isomorphism $F \cong F'$ sends b_i to c'_i and c_i to b'_i for $0 \leq i \leq m+1$. The map $\eta : G_2 \rightarrow \Gamma$ constructed at the beginning of Step 3 is injective on F , so we obtain an epimorphism

$$G_w \rightarrow \Gamma *_F \Gamma'.$$

The latter is an amalgam of virtually free groups along malnormal subgroups, and Wise [42, Theorem 1.3] proved that such amalgams are residually finite. Therefore $\hat{G}_w \not\cong 1$, as required. \square

7. NON-POSITIVELY CURVED SQUARE COMPLEXES

In this section we strengthen Theorem A by proving that the existence of finite-index subgroups remains undecidable among the fundamental groups of compact, non-positively curved square complexes. More precisely, we will prove the geometric form of this result stated in the introduction as Theorem B.

The arguments in this section are topological in nature and the basic construction is close in spirit to earlier constructions by Kan and Thurston [21], Leary [26] and others: the key point in each case is that one replaces a disc in some standard topological construction by a more complicated space that is equally as *inessential* as a disc from one point of view but at the same time admits geometric or topological properties that are more desirable from the point of view of the application at hand. In our setting, the standard construction is that of the 2-complex canonically associated to a group presentation, the desirable property is non-positive curvature, and the appropriate notion of *inessential* is having a profinitely trivial fundamental group, i.e. the spaces that replace the disc should have no connected finite-sheeted coverings.

7.1. An adaptation of the standard 2-complex. Let

$$\mathcal{P} \equiv \langle a_1, \dots, a_n \mid r_1, \dots, r_m \rangle$$

be a finite presentation for a group $G = |\mathcal{P}|$. The standard 2-complex $K(\mathcal{P})$ with fundamental group G is defined as follows: it has a single vertex, a 1-cell for each generator – oriented and labelled a_i – and a 2-cell for each relator, attached along the edge-loop labelled by the word r_j , which we may assume to be cyclically reduced. In what follows, it will be useful to have a name, $R(a_1, \dots, a_n)$ or, more briefly, $R(\underline{a})$, for the 1-skeleton of $K(\mathcal{P})$.

Let X be a compact, non-positively curved square complex with $S = \pi_1 X$ infinite but $\hat{S} \cong 1$ (such as the examples of [13] or [43]) and fix some edge-loop $\gamma : \mathbb{S}^1 \rightarrow X^{(1)}$ in the 1-skeleton that is a local geodesic in X , based at a vertex.

Definition 7.1. Let $S(\mathcal{P})$ be the space obtained by attaching m copies of X to $R(\underline{a})$, with the j -th copy attached by a cylinder joining γ to the edge-loop in $R(\underline{a})$ labelled r_j . More formally, writing $\rho_j : \mathbb{S}^1 \rightarrow R(\underline{a})$ for this last loop, we define \sim to be the equivalence relation on

$$R(\underline{a}) \coprod (\mathbb{S}^1 \times [0, 1]) \times \{1, \dots, m\} \coprod X \times \{1, \dots, m\}$$

defined by

$$\forall t \in \mathbb{S}^1 \forall j \in \{1, \dots, m\} : \rho_j(t) \sim (t, 0, j) \text{ and } (t, 1, j) \sim (\gamma(t), j),$$

and define $S(\mathcal{P})$ to be the quotient space. Define $G_S := \pi_1 S(\mathcal{P})$.

Remarks 7.2. (1) For any fixed choice of γ , the construction of $S(\mathcal{P})$ from \mathcal{P} is algorithmic.

(2) There is a continuous map $\rho : S(\mathcal{P}) \rightarrow K(\mathcal{P})$ that is the identity on $R(\underline{a})$, sends each copy of X to a point in the interior of the corresponding 2-cell of $K(\mathcal{P})$, and maps the interior of each attaching cylinder homeomorphically to the interior of a punctured 2-cell. This map induces epimorphisms $\rho_* : G_S \rightarrow G$ and $\hat{\rho}_* : \hat{G}_S \rightarrow \hat{G}$.

Lemma 7.3. *The map $\hat{\rho}_* : \hat{G}_S \rightarrow \hat{G}$ is an isomorphism.*

Proof. It is enough to show that any homomorphism f from G_S to a finite group factors through ρ_* . By construction, S has no finite quotients, so $f(S_j) = 1$ where $S_j \cong S$ is the fundamental group of the copy of X in $S(\mathcal{P})$ indexed by $j \in \{1, \dots, m\}$. \square

Lemma 7.4. *For any finite group presentation \mathcal{P} , the space $S(\mathcal{P})$ has the structure of a finite, non-positively curved square complex.*

Proof. Let k be the length of γ . We scale $R(a_1, \dots, a_n)$ by a factor of k and subdivide each edge into k pieces of length 1. For $j = 1, \dots, m$ we take a copy of X scaled by a factor of the word-length of r_j , subdivided in the natural way so that it is a (unit) square complex. The attaching maps in the definition of $S(\mathcal{P})$ are then length-preserving, so if the connecting cylinders are subdivided into squares in the obvious manner, $S(\mathcal{P})$ becomes a non-positively curved square complex [10, Proposition II.11.6]. \square

Together, these lemmas establish the following proposition, which reduces Theorem B to Theorem A.

Proposition 7.5. *There is an algorithm that takes as input a finite group presentation \mathcal{P} for a group G and outputs a compact, non-positively curved square complex $S(\mathcal{P})$ with fundamental group G_S such that*

$$\hat{G}_S \cong \hat{G}.$$

Remark 7.6. A simple combinatorial check will determine if a finite square complex satisfies the link condition, i.e. supports a metric of non-positive curvature. Thus, the class of such 2-complexes is recursive.

7.2. Largeness. A group is called *large* (or *as large as a free group*, in the original terminology of Pride [32]), if it has a subgroup of finite index that maps surjectively to a non-abelian free group. Largeness is related to the existence of finite quotients by the following elementary observation.

Lemma 7.7. *A group G has a non-trivial finite quotient if and only if $G * G * G$ is large.*

Proof. If G maps onto a non-trivial finite group Q , then $G * G * G$ maps onto $Q * Q * Q$. The kernel of any homomorphism $Q * Q * Q \rightarrow Q$ that restricts to an isomorphism on each of the free factors is non-abelian and free of finite index, and a subgroup of finite index in $G * G * G$ maps onto it. Conversely, if G can only map trivially to a finite group, then the same is true of $G * G * G$; so it is not large. \square

Combining Lemma 7.7 with Theorem B, we see that largeness is undecidable, even among the fundamental groups of non-positively curved square complexes.

Corollary 7.8. *There is a recursive sequence of finite, non-positively curved square complexes X_n such that:*

- (1) *for each $n \in \mathbb{N}$, X_n has a proper connected finite-sheeted covering space if and only if $\pi_1 X_n$ is large;*
- (2) *the set of natural numbers*

$$\{n \in \mathbb{N} \mid \pi_1 X_n \text{ is large}\}$$

is recursively enumerable but not recursive.

In particular, there is no algorithm to determine whether or not the fundamental group of a finite, non-positively curved square complex is large.

7.3. Biautomatic groups. Fundamental groups of compact, non-positively curved square complexes are biautomatic [18] (see also [28]). There is an algorithm to determine if a biautomatic group is trivial, but Theorem B tells us that there is no algorithm to determine if it is profinitely trivial.

Corollary 7.9. *There is no algorithm that, given a biautomatic group G , can determine whether or not G has a proper subgroup of finite index. Nor is there an algorithm that can determine whether or not G is large.*

8. PROFINITE RANK

By definition, the *profinite rank* of a group G , denoted by $\hat{d}(G)$, is the minimum number of elements needed to generate \hat{G} as a topological group.

8.1. A profinite Grushko lemma. We want to show that there is no algorithm that can determine the profinite rank of a hyperbolic group. For this we shall use the following analogue of Grushko's theorem; we make no claim that the constant $\frac{59}{60}$ is sharp.

Lemma 8.1. *Let G be a group with $\hat{G} \not\cong 1$. Then $\hat{d}(*_{i=1}^n G) \geq \frac{59}{60}n$.*

Proof. If G maps onto a finite cyclic group \mathbb{Z}/p , then $L_n := *_{i=1}^n G$ and its profinite completion map onto $(\mathbb{Z}/p)^n$, and therefore require at least n generators.

Suppose, then, that G maps onto a non-trivial finite perfect group S . Let $Q_n := *_{i=1}^n S$ and let $\pi : Q_n \rightarrow S$ be a homomorphism that restricts to an isomorphism on each free factor. The kernel $\ker \pi$ acts freely on the Bass–Serre tree for Q_n (since all of the torsion of Q_n is conjugate into one of the free factors) and hence $\ker \pi$ is a free group; its rank is $r := (n-1)(|S|-1)$, as can be calculated using rational Euler characteristic.

Let $K_n < L_n$ be the inverse image of $\ker \pi$. Then K_n is normal, maps onto a free group of rank r , and $L_n/K_n \cong S$. We fix an epimorphism $K_n \rightarrow (\mathbb{Z}/2)^r =: A$ and induce this to a homomorphism $\Phi : L_n \rightarrow A \wr S$. The image of K_n under this map lies in the base of the wreath product, where it projects onto each A summand; thus it is an elementary 2-group of rank at least r .

By the Nielsen–Schreier formula, if $\Phi(L_n)$ has rank δ then $\Phi(K_n)$, which has index at most $|S|$, has rank at most $|S|(\delta-1)+1$. Thus

$$(n-1)(|S|-1) \leq |S|(\delta-1)+1,$$

whence

$$\hat{d}(L_n) \geq \delta \geq \left(\frac{|S|-1}{|S|} \right) n.$$

But S is perfect and non-trivial, so $|S| \geq 60$. □

8.2. Profinite rank of hyperbolic groups. We shall appeal to the following version of the Rips construction.

Theorem 8.2. *There is an algorithm that takes as input a finite presentation for a group G and outputs a finite presentation for a residually finite, torsion-free, hyperbolic group Γ such that there exists a short*

exact sequence

$$1 \rightarrow N \rightarrow \Gamma \rightarrow G \rightarrow 1$$

where N is a 2-generator group.

Proof. Rips showed how to construct such a short exact sequence with Γ satisfying the $C'(1/6)$ small-cancellation condition [34]. Wise proved that such groups are fundamental groups of compact, non-positively curved cube complexes [39]. By Agol's theorem [3], it follows that Γ is virtually special and, in particular, residually finite. \square

We can now prove part (5) of Theorem D. Note that the examples constructed are residually finite.

Theorem 8.3. *Fix any $d_0 > 2$. There is a recursive sequence of torsion-free, residually finite, hyperbolic groups Γ_n with the property that:*

- (1) *for any $n \in \mathbb{N}$, $\hat{d}(\Gamma_n) < d_0 \Leftrightarrow \hat{d}(\Gamma_n) = 2$; and*
- (2) *the set of natural numbers*

$$\{n \in \mathbb{N} \mid \hat{d}(\Gamma_n) \geq d_0\}$$

is recursively enumerable but not recursive.

In particular, there is no algorithm that can decide whether or not the profinite completion of a torsion-free, residually finite, hyperbolic group can be generated (topologically) by a set of cardinality less than d_0 .

Proof. Let G_n be a recursive sequence of finitely presented groups such that the set of natural numbers $\{n \in \mathbb{N} \mid \hat{G}_n \not\cong 1\}$ is recursively enumerable but not recursive. Let $M \geq \frac{60}{59}d_0$ and, for each n , let G'_n be a free product of M copies of G_n . Then either $\hat{G}'_n \cong 1$ or $\hat{d}(G'_n) \geq d_0$ by Lemma 8.1.

Apply Theorem 8.2 to obtain short exact sequences

$$1 \rightarrow N_n \rightarrow \Gamma_n \rightarrow G'_n \rightarrow 1$$

with each N_n a 2-generator group. Note that since Γ_n is residually-finite but not cyclic, $\hat{d}(\Gamma_n) \geq 2$.

If $\hat{d}(\Gamma_n) < d_0$ then $\hat{d}(G'_n) < d_0$, so $\hat{G}'_n \cong 1$ and \hat{N}_n surjects $\hat{\Gamma}_n$, whence $\hat{d}(\Gamma_n) = 2$. This proves (1). Item (2) follows, because $\hat{d}(\Gamma_n) \geq d_0$ if and only if $\hat{G}_n \not\cong 1$. \square

9. UNDECIDABLE PROPERTIES OF HYPERBOLIC GROUPS

In this section we prove the remaining parts of Theorem D. We also prove that either every hyperbolic group is residually finite, or else there is no algorithm to decide which hyperbolic groups have a finite quotient.

All of these things will be proved by combining our previous results with the following refinement of the Rips construction [34], which is due to Belegradek and Osin [6].

Theorem 9.1 (Belegradek–Osin, [6]). *There is an algorithm that takes as input a finite presentation for a non-elementary hyperbolic group H and a finite presentation for a group G and outputs a presentation for a hyperbolic group Γ that fits into a short exact sequence*

$$1 \rightarrow N \rightarrow \Gamma \rightarrow G \rightarrow 1$$

such that N is isomorphic to a quotient group of H . Furthermore, if H and G are torsion-free then Γ can also be taken to be torsion-free.

Proof. The only point that is not addressed directly by Belegradek and Osin is the fact that the construction can be made algorithmic, but it is tacitly implied in Corollary 3.8 of [6]. Indeed, since the class of hyperbolic groups is recursively enumerable [31], a naive search will eventually find a hyperbolic group Γ and a homomorphism $H \rightarrow \Gamma$ whose image is normal with quotient isomorphic to G .

In the torsion-free case, one needs the well known fact that the class of torsion-free hyperbolic groups is also recursively enumerable (see, for instance, the proof of Theorem III.Γ.3.2 in [10]). \square

9.1. Largeness and virtual first Betti number. Parts (1) and (2) of Theorem D follow from the next theorem.

Theorem 9.2. *There is a recursive sequence of finite presentations for torsion-free, hyperbolic groups Γ_n such that:*

(1) *for each $n \in \mathbb{N}$,*

$$vb_1(\Gamma_n) > 0 \Leftrightarrow vb_1(\Gamma_n) = \infty \Leftrightarrow \Gamma_n \text{ is large ;}$$

and

(2) *the set of natural numbers*

$$\{n \in \mathbb{N} \mid \Gamma_n \text{ is large}\}$$

is recursively enumerable but not recursive.

In particular, for any $1 \leq d \leq \infty$, there is no algorithm that determines whether or not a given torsion-free hyperbolic group Γ has $vb_1(\Gamma) \geq d$; likewise, there is no algorithm that determines whether or not a given torsion-free hyperbolic group is large.

Proof. Let G_n be the sequence of fundamental groups of the square complexes produced by Corollary 7.8; note that as the fundamental groups of aspherical spaces, the G_n are torsion-free. Let $N_n < \Gamma_n$ be the pair of groups obtained by applying the algorithm of Theorem 9.1 to

G_n , with H a fixed torsion-free, non-elementary hyperbolic group with Property (T); torsion-free uniform lattices in $\mathrm{Sp}(n, 1)$ provide explicit examples.

We have the following chain of implications.

$$vb_1(G_n) > 0 \Rightarrow \Gamma_n \text{ is large} \Rightarrow vb_1(\Gamma_n) = \infty \Rightarrow vb_1(\Gamma_n) > 0$$

The first implication follows from part (1) of Corollary 7.8, and the other implications are trivial.

To prove (1) and (2), it therefore suffices to show that $vb_1(\Gamma_n) > 0$ implies that $vb_1(G_n) > 0$. Suppose, therefore, that $K < \Gamma_n$ is a subgroup of finite index that admits a surjection $f : K \rightarrow \mathbb{Z}$. Property (T) is inherited by quotients and subgroups of finite index, so the abelianization of $N_n \cap K$ is finite. Therefore, $f(N_n \cap K) = 1$ and so $K/(K \cap N_n)$ surjects \mathbb{Z} . But $K/(K \cap N_n)$ has finite index in G_n , so $vb_1(G_n) > 0$ as required. \square

9.2. Linear representations. In this section we make use of known examples of torsion-free, non-elementary hyperbolic groups that admit no infinite linear representation to establish parts (3) and (4) of Theorem D. As M. Kapovich showed in [23, Theorem 8.1], the existence of such examples can be proved using the work of Corlette [16] and Gromov–Schoen [19] on (archimedean and non-archimedean) super-rigidity for lattices in $\mathrm{Sp}(n, 1)$.

Theorem 9.3 ([23]). *There exists a torsion-free, non-elementary hyperbolic group H with the property that, for any field k , every finite-dimensional representation of G over k has finite image.*

Proof. The statement of this theorem is the same as [23, Theorem 8.1], with the additional stipulation that the group H is torsion-free. Following Kapovich, we start with a uniform lattice Γ in the isometry group of quaternionic hyperbolic space. By Selberg’s Lemma, we may assume that Γ is torsion free. We then take H (which is G in Kapovich’s notation) to be any infinite small-cancellation quotient of Γ . As Kapovich explains, the group H then has no infinite linear representations over any field.

For a suitable choice of small-cancellation quotient, any torsion in H is the image of torsion in Γ . (For instance, this follows from [30, Lemma 6.3], which even deals with the relatively hyperbolic setting.) Such a choice of H is therefore torsion-free. \square

The following theorem covers parts (3) and (4) of Theorem D.

Theorem 9.4. *Fix any infinite field k . There is a recursive sequence of torsion-free hyperbolic groups Γ_n with the property that:*

- (1) *for any $n \in \mathbb{N}$, Γ_n has a finite-dimensional representation over k with infinite image if and only if Γ_n has a finite-dimensional representation over some field with infinite image; and*
- (2) *the set of $n \in \mathbb{N}$ such that Γ_n has a finite-dimensional representation over k with infinite image is recursively enumerable but not recursive.*

Proof. Let X_n be the sequence of square complexes output by Corollary 7.8 and let $G_n = \pi_1 X_n$. Finitely generated linear groups are residually finite, so for any infinite field k , G_n has a finite-dimensional representation over k with infinite image if and only if G_n is large; furthermore, the set of natural numbers n such that G_n has such a representation is recursively enumerable but not recursive.

Let H be the torsion-free, non-elementary hyperbolic group of Theorem 9.3. For each n , let Γ_n be the torsion-free hyperbolic group that is the output of the algorithm of Theorem 9.1 with input G_n and H .

The result now follows from the claim that, for any field k , Γ_n has a finite-dimensional representation over k with infinite image if and only if G_n does. Indeed, if G_n has such a representation then Γ_n clearly does. Conversely, suppose that $f : \Gamma_n \rightarrow \mathrm{GL}(m, k)$ has infinite image. If N is the kernel of the map $\Gamma_n \rightarrow G_n$ then, because N is a quotient of H , it follows that $f(N)$ is finite. Because $f(\Gamma_n)$ is residually finite, there exists a proper subgroup K of finite index in $f(\Gamma_n)$ such that $K \cap f(N) = 1$. Then $L = f^{-1}(K)$ is a subgroup of finite index in Γ_n with an infinite representation $f|_L$ over k , and $f|_L(L \cap N) = 1$. Therefore, $f|_L$ factors through the restriction to L of the map $\Gamma_n \rightarrow G_n$. It follows that G_n has a subgroup of finite index with an infinite representation over k , and so G_n also has such a representation. \square

9.3. Profinite undecidability in the hyperbolic case. We finish with the following conjecture.

Conjecture 9.5. *There is no algorithm that can determine whether or not a given hyperbolic group Γ has $\hat{\Gamma} \cong 1$.*

Since the triviality problem is solvable for hyperbolic groups, the above conjecture is false if every non-trivial hyperbolic group Γ has $\hat{\Gamma} \not\cong 1$. In fact, I. Kapovich and Wise proved that every non-trivial (torsion-free) hyperbolic group Γ has $\hat{\Gamma} \not\cong 1$ if and only if every (torsion-free) hyperbolic group is residually finite [22]. Conjecture 9.5 therefore implies the well known conjecture that there exists a non-residually finite hyperbolic group [7, Question 1.15]. In fact, our final theorem shows that the two conjectures are equivalent (even in the torsion-free case).

Theorem 9.6. *The following statements are equivalent.*

- (1) *Every non-trivial (torsion-free) hyperbolic group has a proper subgroup of finite index.*
- (2) *There is an algorithm that, given a finite presentation of a (torsion-free) hyperbolic group, will determine whether or not the profinite completion of that group is trivial.*

Proof. There is an algorithm that can decide if a given hyperbolic group is trivial, and if (1) holds then (2) reduces to checking if the given group is trivial. For the converse, suppose that there exists a non-trivial hyperbolic group H_0 with $\hat{H}_0 = 1$. Clearly H_0 is non-elementary. Let G_n be a sequence of (torsion-free) groups that witnesses the undecidability in Theorem B, let Γ_n be the sequence of hyperbolic groups obtained by applying Theorem 9.1 to G_n with $H = H_0$, and note that $\hat{H} = 1$ implies $\hat{\Gamma}_n \cong \hat{G}_n$. It is a feature of Theorem 9.1 that if H_0 is torsion-free then so are the groups Γ_n . \square

REFERENCES

- [1] S. I. Adyan. Algorithmic unsolvability of problems of recognition of certain properties of groups. *Dokl. Akad. Nauk SSSR (N.S.)*, 103:533–535, 1955.
- [2] S. I. Adyan. Unsolvability of some algorithmic problems in the theory of groups. *Trudy Moskov. Mat. Obšč.*, 6:231–298, 1957.
- [3] Ian Agol. The virtual Haken conjecture. *Documenta Math.*, 18:1045–1087, 2013, with an appendix by Ian Agol, Daniel Groves and Jason Manning.
- [4] Jitendra Bajpai. Omnipotence of surface groups. *Masters Thesis, McGill University*, 2007.
- [5] Gilbert Baumslag, W. W. Boone, and B. H. Neumann. Some unsolvable problems about elements and subgroups of groups. *Math. Scand.*, 7:191–201, 1959.
- [6] Igor Belegradek and Denis Osin. Rips construction and Kazhdan property (T). *Groups Geom. Dyn.*, 2(1):1–12, 2008.
- [7] Mladen Bestvina. Questions in geometric group theory. <http://www.math.utah.edu/~bestvina/eprints/questions-updated.pdf>.
- [8] Meenaxi Bhattacharjee. Constructing finitely presented infinite nearly simple groups. *Comm. Algebra*, 22(11):4561–4589, 1994.
- [9] William W. Boone. The word problem. *Ann. of Math. (2)*, 70:207–265, 1959.
- [10] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [11] Martin R. Bridson and Henry Wilton. The isomorphism problem for profinite completions of finitely presented, residually finite groups. *Groups Geom. Dyn.*, to appear, 2014.
- [12] Martin R. Bridson and Henry Wilton. Undecidability and the developability of permutoids and rigid pseudogroups. *arXiv:1405.4368*, 2014.
- [13] Marc Burger and Shahar Mozes. Finitely presented simple groups and products of trees. *C. R. Acad. Sci. Paris Sér. I Math.*, 324(7):747–752, 1997.

- [14] Jack O. Button. Largeness of LERF and 1-relator groups. *Groups Geom. Dyn.*, 4(4):709–738, 2010.
- [15] Peter Cameron. Extending partial permutations. <http://www.maths.qmul.ac.uk/~pjc/odds/partial.pdf>, 2004.
- [16] Kevin Corlette. Archimedean superrigidity and hyperbolic geometry. *Ann. of Math. (2)*, 135(1):165–182, 1992.
- [17] M. Dehn. Über unendliche diskontinuierliche Gruppen. *Math. Ann.*, 71(1):116–144, 1911.
- [18] S. M. Gersten and H. B. Short. Small cancellation theory and automatic groups. *Invent. Math.*, 102(2):305–334, 1990.
- [19] Mikhail Gromov and Richard Schoen. Harmonic maps into singular spaces and p -adic superrigidity for lattices in groups of rank one. *Inst. Hautes Études Sci. Publ. Math.*, 76:165–246, 1992.
- [20] Charles F. Miller III. Decision problems for groups—survey and reflections. In *Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989)*, number 23 in Math. Sci. Res. Inst. Publ., pages 1—59. Springer, New York, 1992.
- [21] D. M. Kan and W. P. Thurston. Every connected space has the homology of a $K(\pi, 1)$. *Topology*, 15(3):253–258, 1976.
- [22] Ilya Kapovich and Daniel T. Wise. The equivalence of some residual properties of word-hyperbolic groups. *J. Algebra*, 223(2):562–583, 2000.
- [23] Michael Kapovich. Representations of polygons of finite groups. *Geom. Topol.*, 9:1915–1951 (electronic), 2005.
- [24] O. G. Kharlampovich. The universal theory of the class of finite nilpotent groups is undecidable. *Mat. Zametki*, 33(4):499–516, 1983.
- [25] Olga Kharlampovich and Alexei Myasnikov. Decidability of the elementary theory of a torsion-free hyperbolic group. *arXiv:1303.0760v4*, 2013.
- [26] Ian J. Leary. A metric Kan-Thurston theorem. *J. Topol.*, 6(1):251–284, 2013.
- [27] G. S. Makanin. Decidability of the universal and positive theories of a free group. *Izvestiya Akademii Nauk SSSR. Seriya Matematicheskaya*, 48(4):735–749, 1984.
- [28] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998.
- [29] P. S. Novikov. *On the algorithmic unsolvability of the word problem in group theory*. Trudy Mat. Inst. im. Steklov. no. 44. Izdat. Akad. Nauk SSSR, Moscow, 1955.
- [30] Denis Osin. Small cancellations over relatively hyperbolic groups and embedding theorems. *Ann. of Math. (2)*, 172(1):1–39, 2010.
- [31] P. Papasoglu. An algorithm detecting hyperbolicity. In *Geometric and computational perspectives on infinite groups (Minneapolis, MN and New Brunswick, NJ, 1994)*, volume 25 of DIMACS Ser. Discrete Math. Theoret. Comput. Sci., pages 193–200. Amer. Math. Soc., Providence, RI, 1996.
- [32] Stephen J. Pride. The concept of “largeness” in group theory. In *Word problems, II (Conf. on Decision Problems in Algebra, Oxford, 1976)*, volume 95 of Stud. Logic Foundations Math., pages 299–335. North-Holland, Amsterdam, 1980.
- [33] Michael O. Rabin. Recursive unsolvability of group theoretic problems. *Ann. of Math. (2)*, 67:172–194, 1958.

- [34] E. Rips. Subgroups of small cancellation groups. *The Bulletin of the London Mathematical Society*, 14(1):45–47, 1982.
- [35] Z. Sela. Diophantine geometry over groups. VII. The elementary theory of a hyperbolic group. *Proc. Lond. Math. Soc. (3)*, 99(1):217–273, 2009.
- [36] A. M. Slobodskoi. Undecidability of the universal theory of finite groups. *Algebra i Logika*, 20(2):207–230, 251, 1981.
- [37] John R. Stallings. Topology of finite graphs. *Inventiones Mathematicae*, 71(3):551–565, 1983.
- [38] Henry Wilton. Virtual retractions, conjugacy separability and omnipotence. *J. Algebra*, 323:323–335, 2010.
- [39] D. T. Wise. Cubulating small cancellation groups. *Geom. Funct. Anal.*, 14(1):150–214, 2004.
- [40] D. T. Wise. The structure of groups with a quasi-convex hierarchy. Preprint, April 2012.
- [41] Daniel T. Wise. Subgroup separability of graphs of free groups with cyclic edge groups. *The Quarterly Journal of Mathematics*, 51(1):107–129, 2000.
- [42] Daniel T. Wise. The residual finiteness of negatively curved polygons of finite groups. *Inventiones Mathematicae*, 149(3):579–617, 2002.
- [43] Daniel T. Wise. Complete square complexes. *Comment. Math. Helv.*, 82(4):683–724, 2007.

MATHEMATICAL INSTITUTE, ANDREWS WILES BUILDING, RADCLIFFE OBSERVATORY QUARTER, OXFORD OX2 6GG, UK

E-mail address: `bridson@maths.ox.ac.uk`

DEPARTMENT OF MATHEMATICS, UNIVERSITY COLLEGE LONDON, GOWER STREET, LONDON WC1E 6BT, UK

Current address: DPMMS, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WB, UK

E-mail address: `h.wilton@maths.cam.ac.uk`